A Normative Account of the Need for Explanation

Wai-hung Wong & Zanja Yudell

Although explanation is a central topic in the philosophy of science and several important models of scientific explanation have been developed and widely discussed, there is one issue concerning explanation that has drawn very little attention, namely, why some phenomena[1] need explanation while some do not. Some may think that any phenomenon that has not been explained needs explanation. Others may think that if we divide (heretofore unexplained) phenomena into those that need explanation and those that do not, it is merely because we are puzzled by, or curious about, the former phenomena but not the latter. According to this line of thinking, the question of why some phenomena need explanation while some do not is a psychological question in disguise: the question is in fact why some people are puzzled [2] by some, but not all, phenomena and want to have an explanation of them. In this paper we will argue that the distinction between phenomena that need explanation and those that do not need one is philosophical, not psychological, because the distinction can be drawn normatively. We will propose a *normative* account of the need for explanation.

In section I, we will explain the difference between a psychological account and a normative account of the need for explanation, and why a normative account is philosophically significant. Section II will be a criticism of what we call *the surprise account* of the need for explanation, which seems more like a psychological account than a normative account. We will also explain why the account is appealing and partially successful. In section III we will discuss a more recent account suggested by Stephen R. Grimm (2008), which can be called *the fact-and-foil account*. It is not clear whether the fact-and-foil account is meant to be psychological or normative; we will argue that as a normative account it is clearly inadequate. Section IV will be a development of our own account of the need for explanation, which we call *the map account*. In section V we will explain why the map account is normative. We will end the paper with a few concluding remarks in section VI.

I. Why an Account of the Need for Explanation?

When something has happened and we know why it happened, we don't typically ask for an explanation for the simple reason that we already have one.[3] It is not the case, however, that *whenever* something has happened and we don't know why it happened, we ask for an explanation. Sometimes we decide not to ask for an explanation because of practical considerations, but in many cases we don't ask for an explanation simply because we don't think what has happened *needs an*

[1] Following Hempel, we use the term 'phenomenon' to refer to the explanandum. See Hempel (1965: 247).

[2] To avoid clumsiness, from now on we will use 'puzzled by' to mean 'puzzled by or curious about' and use 'puzzlement' to refer to both puzzlement and curiosity.

[3] Sometimes we don't ask for an explanation merely because *we believe*, rightly or wrongly, that we know why it happened.

explanation. Thinking that something does not need an explanation is not the same as thinking that it does not have an explanation. Even a person who believes that everything has an explanation may not think everything needs an explanation. If, for example, an orange grows to the size of a watermelon, we certainly would think an explanation is required; but if an orange grows to the size of a grapefruit, we would not think so. It is not just in the case of natural phenomena that we may not think something that has happened needs an explanation even when we don't have one. If a colleague were to win the lottery, we would not think his winning needs explanation; but if he won the lottery three times in a row and every time it was exactly one day after he received a fortune cookie saying that he would win the lottery the next day, then we would think that his winning in this way requires an explanation.

But why is there such a difference? Why is it not the case that *whenever* we don't know why something has happened, we think that the phenomenon in question needs explanation? Some may suggest that the difference is simply a psychological one: we are puzzled by some phenomena but not by others; if we are not puzzled, then it is just natural that we don't think an explanation is needed. Now if what it is to be puzzled by a phenomenon is to think that an explanation of the phenomenon is needed, then it is trivially true that we think a phenomenon needs explanation if and only if we are puzzled by it. However, puzzlement can be understood as a psychological state such that (1) when S is puzzled by P, S's puzzlement can be identified independently of S's thought that P needs explanation, and (2) S's puzzlement can cause S's thought that P needs explanation. On the view suggested, it is usually, if not always, the psychological state of puzzlement that causes the thought that a particular phenomenon needs explanation.

If this view is correct, there may not be anything philosophically significant about the need for explanation — different phenomena make different people feel puzzled because these people have different psychological backgrounds, and the need for explanation, which is caused by such puzzlement, is in this sense subjective. On this view, even if there are patterns or regularities concerning when people ask for an explanation of a phenomenon, those are just psychological patterns or regularities that reveal nothing about the nature of explanation. If, for example, a certain group of people (which can be a very large group) in circumstances C all (or mostly) think phenomenon P needs explanation, it is because they share the same psychological background and hence are puzzled by P in C. Such a pattern, however, does not tell us anything about how the explanation of P can be related to the explanation of other phenomena, whether there are any criteria for judging whether an explanation of P is successful, or how a successful explanation of P is related to our broader understanding of the world. These people feel the need for an explanation of P; some of them may try to find one; and if a successful explanation is found, they will be happy. That's all.

It is worth exploring another kind of account of the need for explanation, that is, a normative one rather than a psychological one. On a normative account, a phenomenon needs an explanation only when we have a good reason to ask for an explanation of it. It is not because we want to have an explanation of a phenomenon that it needs an explanation; it is because the phenomenon needs an explanation that we *should* ask for an explanation of it. Although the concept of normativity is notoriously difficult to pin down and there are different theories of normativity, it should be clear from what we have just said that we are using 'normative' mainly in the directive sense rather than the evaluative sense. A normative statement in this sense is about action, such as

'You should look for an explanation of this phenomenon', and such a statement is true when there are reasons for the action involved.[4]

Our account does not, however, rely on any specific theory of normativity. And we will leave the notion of normativity (in the directive sense) flexible: normative statements, as Christine Korsgaard puts it, "*in various ways*, claim to direct us, to guide our thoughts, desires, and actions" (Korsgaard 1996: 20; italics added); they can "command, oblige, recommend, or guide" (8). Accordingly, if someone agrees that a phenomenon needs an explanation but fails to look for one (even when she has the time and resources to do so), she is being unreasonable or irrational to some extent. To what extent? It depends on the case. In some cases we may have to say she has violated an epistemic obligation. In these cases, the relation between 'S has a good reason to ask for an explanation of P' and 'S ought to look for an explanation of P' is analogous to that between 'S has sufficient evidence for P' and 'S ought to believe that P'. If S has sufficient evidence for P, then her believing that P is not only permitted, but also required. If she does not believe that P (or even believe that not-P), she does something wrong epistemically. (What these cases may be will become clearer after we explain our account of the need for explanation.)

Now there could be an account of the need for explanation according to which the very fact that a phenomenon has not been explained is a good reason to ask for an explanation of the phenomenon. Such an account is in a sense normative, but this would not be an account of why some phenomena need explanation while some do not, for on such an account every heretofore unexplained phenomenon needs explanation — there is not really a substantive distinction between phenomena that need explanation and those that do not. The normative account we will be suggesting is precisely an attempt to draw a distinction between phenomena that need explanation and those that do not. According to our account, not every heretofore unexplained phenomenon is such that we have a good reason to ask for an explanation of it; only those for which we have reason to ask for an explanation need an explanation.

In the remainder of this section we will explain why it is philosophically significant to have a normative account of the need for explanation. First of all, such an account can shed light on the nature of explanation. Philosophical discussions concerning explanation have been focusing on the structure of *scientific* explanation and the criteria of success for such explanation, but as James Woodward remarks,

> it is implicitly assumed in most discussions of scientific explanation that there are important similarities or continuities in structure between explanations like (2.4.1) and explanations that are more obviously scientific and that these similarities that should be captured by some common account that applies to both. Indeed, it is a striking feature not just of Hempel (1965) but of many other treatments of scientific explanation that much of the discussion in fact focuses on "ordinary life" singular causal explanations similar to (2.4.1), the tacit assumption being that conclusions about the structure of such explanations have fairly direct implications for understanding explanation in science. (Woodward 2009)

[4] For a clear exposition of the distinction between directives and evaluatives, see Thomson (2008). As Thomson points out, many philosophers "think that what makes directives true, when they are, is facts about reasons for action" (125); we agree with these philosophers. Thomson argues, however, that "directives do not rest on reasons for action, rather reasons for action rest on directives" (126).

The example (2.4.1) in Woodward's article is "the impact of my knee on the desk caused the tipping over of the inkwell". Let us call explanations like (2.4.1) *everyday explanations*. The common account Woodward mentions is a common account of the structure of both scientific explanation and everyday explanation. Such a common account assumes that scientific explanation and everyday explanation are similar in structure in important respects, and the account is supposed to capture the similarities. But what reason do philosophers of science have for assuming that there are such structural similarities between scientific explanation and everyday explanation? A normative account of the need for explanation can help answer this question: if the need for explanation arises in the same way in both science and everyday life as a *rational* response to a phenomenon that has not been explained, it is reasonable to assume that what would satisfy the need for explanation in science and what would satisfy the need for explanation in everyday life have similar structural features, for it is reasonable to assume that they are both subject to similar rational constraints or requirements.

Why can't, some may ask, a psychological account of the need for explanation provide us with some reason for assuming that there are structural similarities between scientific explanation and everyday explanation? Here is a brief answer: if it were *merely* some kind of psychological need that motivated us to look for an explanation of a phenomenon, we would have no reason to believe that an explanation that would satisfy the psychological need has to have certain structural features; the psychological need could be satisfied rather differently in the case of different people or groups of people even under similar circumstances. Perhaps a psychological account of the need for explanation coupled with an ingenious psychological theory might help show that scientific explanation and everyday explanation have structural similarities, but until we find such a psychological theory, a normative account is clearly preferable.

A normative account of the need for explanation can also help us understand better the development and progress of science. Although the development and progress of science depends on many factors, some of which are psychological or sociological, the need for explanation should be considered a determining factor. It is when there is a need for an explanation of a phenomenon that an explanation of it is sought, and the search for an explanation of one phenomenon rather than another pushes the science involved to go one direction rather than another. Unlike psychological and sociological factors, the need for explanation, understood in the normative way we have suggested, is constrained by some rational requirements. To the extent to which the development and progress of science is determined by the need for explanation, it can be understood as having a direction that is rational, whether the aim of science is, as scientific realists think, to give us more and more truths about the world, or, as constructive empiricists think, to give us theories that are more and more empirically adequate.

II. The Surprise Account

There is a common, intuitive answer to the question of when an explanation is needed: the need for explanation of a phenomenon is connected with how surprising or unexpected the phenomenon is. Larry Laudan, for example, seems to be expressing such a view when he remarks that "anything about the natural world which strikes us as *odd*, or otherwise in need of explanation, constitutes an empirical problem [which scientific theories are designed to solve]" (Laudan 1977: 15; italics added). Although he adds the qualifying clause 'or otherwise in need of explanation', the main idea

is that a phenomenon that we think is in need of explanation is in most cases one that we find surprising or unexpected. This answer may not have been proposed as a formal theory of the need for explanation, but it will be worth considering it as such. Thus, let us take the *surprise account* of the need for explanation to consist of the following main claim: a phenomenon needs explanation to the degree that the phenomenon is surprising or unexpected.[5]

As stated, the surprise account is ambiguous. On the one hand, we may consider the need for explanation on this account to be a mere *psychological* need. On such an interpretation, the account is attempting to capture when, as a matter of fact, we tend to be puzzled by a phenomenon — we are puzzled when the phenomenon is surprising — and as a result think that it needs explanation. Such an account would be successful to the degree that it predicts when people actually feel a need for explanation. A proper evaluation of such an account would presumably involve empirical studies of large numbers of people. On the other hand, we may consider the need for explanation to be a *normative* need even on this account. On this interpretation, the account is trying to make sense of when we have good reason to require explanation — that a phenomenon is surprising is a good reason to require explanation of it. The measure of success for such an account will be less closely connected to when we actually do make or feel the need for explanation. It might be the case that there are circumstances in which it is common to require explanation where the need for explanation is not warranted, or circumstances in which we typically do not require explanation although we ought to do so. So, the success of such an account should be measured by how well it captures the normative character of the need for explanation.

However, since finding something surprising, like being puzzled, is psychological in nature, and since it is not clear how finding a phenomenon surprising *in itself* can be a good reason to require an explanation of the phenomenon, the surprise account looks more like a psychological account than a normative account. In any case, we do not need to determine whether the surprise account is a psychological or normative account, for it will not matter to our criticism of it. Whether taken as normative or psychological, the surprise account makes the surprisingness of a phenomenon both necessary and sufficient for explanation to be needed of the phenomenon. As it turns out, surprisingness is neither necessary nor sufficient for explanation to be needed, and it is instructive to see why.

As Stephen Grimm (2008: 485) points out, there are cases in which a phenomenon is not surprising, yet explanation is needed; in these cases, surprisingness is not necessary for explanation to be needed. For example, the squeaking of a bicycle wheel may be quite familiar to and hence expected by its owner, but for him it still needs explanation. Such examples may be multiplied quite easily and found in many domains: the rising of the tide, the course of a disease, and the acquisition of language by a toddler are all phenomena that we expect but stand in need of explanation (or once stood in need of explanation). In all these cases, the phenomenon belongs to a type that we have come to expect through inductive inference, which explains why the phenomenon is not surprising; the familiarity does not, however, make the phenomenon any less in need of explanation.

Grimm considers a modification to the surprise account that attempts to avoid this kind of difficulty. Let the account say that a phenomenon needs explanation to the degree that the phenomenon is *ur*-surprising, where ur-surprisingness measures "how likely one would have taken

[5] For simplicity, we will be using the word 'surprising' instead of the phrase 'surprising or unexpected'.

the situation to be *before* one first learned of the situation" (485, footnote 6). Given that every rising of the tide or occurrence of cancer seems to us a new situation, we don't think this solution goes very far, but if we understand 'situation' to mean 'type of phenomenon', then ur-surprisingness promises a more robust alternative account. Nonetheless, as Grimm points out, there are phenomena and types of phenomena that are *never* surprising yet seem to need explanation. He gives the example of the greenness of tree leaves, which may only come to demand an explanation once one learns that leaves may change color. In this case, the greenness of leaves was never surprising, and it is learning about some other phenomenon that makes the greenness of the leaves in need of explanation. Again, we may easily multiply examples.

Indeed, there are cases where a phenomenon is fully expected and yet needs explanation on its first appearance. For example, when astronomers first turned their telescopes to the region determined by Le Verrier's calculations from the irregularities in the orbit of Uranus, they were not surprised to discover Neptune there. The question of why Neptune was there arose as soon as it was discovered, and was not made in need of explanation by any subsequent knowledge. In general, as Wesley Salmon (1978) observes, we can be led to the discovery of a cause by knowing its effect without thereby having an explanation for the cause. In many such cases the need for explanation of the cause will arise immediately upon learning of the cause, although the cause itself is not a surprise.

This last kind of case suggests a modification to the ur-surprisingness account that requires making the account explicitly involve normative concepts. Consider the case of the discovery of Neptune. One might be tempted to say that the discovery of Neptune should have been made when the irregularities in the orbit of Uranus were first discovered. Although it took some time for the implication of those irregularities to be recognized, that delay was a result of our flawed cognitive capacities, and we should have known everything there was to know about the orbit of Neptune then without needing empirical confirmation. Thus, according to this thought, the need for the explanation of Neptune's existence corresponds to how surprising it *would have been* when it *should have been* discovered. The idea can be generalized by saying that ur-surprisingness is a measure of how likely one would have taken the situation to be *before* one first should have learned of the situation.

Now setting aside any difficulties with this conception of how scientific discovery works, it seems evident that this more normative version of the surprise account avoids the problem introduced by the knowledge that effects give of causes. If the discovery of an effect is tantamount to the discovery of the cause, then prior to the discovery of the effect the cause is unexpected and hence in need of explanation. However, this normative conception does not help with Grimm's original objection that a known phenomenon may come to need explanation as a result of subsequent discoveries, though it was never surprising. Nonetheless, this explicit invocation of normativity points the way toward a better account of the need for explanation, and our view will be explicitly normative in much this way.

Although it seems clear that the surprise account fails by virtue of the fact that surprisingness is not *necessary* for the need for explanation, it's worth examining how it fails in the other direction as well. Surprisingness is not *sufficient* for explanation to be needed either. Consider a typical lottery winner. It is not surprising at all that someone wins the lottery. However, the lottery winner sure is (or should be anyway) surprised that she won the lottery! Here is a paradigmatic example of a surprising occurrence, yet it is clear that no explanation is typically

needed. Indeed, this example also highlights the normative character of the need for explanation, for it is also a paradigmatic case of when someone looks for explanation when she ought not do so because there is no good reason to ask for explanation of the phenomenon. There are surely many lottery winners who do seek for explanation of their good fortune (Was it because I chose my children's birthday numbers? Was it a sign from God?) that typically[6] needs no explanation.

We think that this counterexample clearly shows the insufficiency of the surprise account. It also raises the important question of why some such surprising phenomena stand in need of explanation and others do not. In the case of the lottery, it seems that the answer lies in the fact that we have reason to think of the outcome of the lottery as a *chancy* [7] event, and thus inexplicable. We will re-characterize this case in section IV according to the account of the need for explanation we propose.

Although it seems that we have driven the surprise account into the ground, a lot has been learned from its failure. There is an important insight in the surprise account that we attempt to incorporate into our proposed account. Properly qualified, the surprisingness of a phenomenon is indeed relevant to its need for explanation. But the failure of the basic surprise account shows that the surprisingness must be relative to something other than our actual state of expectation. The failure of the ur-surprisingness account shows that we cannot simply relativize surprisingness to our state of expectation at a time. We further gave some grounds for thinking that the account must be explicitly normative. Finally, the insufficiency of the surprise account also suggests that the properly relativized notion of surprisingness must not only draw in some actually unsurprising phenomena, it must also exclude some phenomena that are in fact surprising.

We draw on all these insights in building our account in section IV. However, before doing so, we turn our attention to a more recent proposal for accounting for the need for explanation proposed by Grimm (2008). We believe that his proposal fails as well, but its failure also provides further insight into how to construct a better proposal.

III. The Fact-and-Foil Account

Recently, Grimm (2008) has offered an account that directly addresses the issue of when an explanation is needed. Grimm says that his account, which we will call the *fact-and-foil account*, is meant to answer the question "What is it *in virtue of which* a situation stands in need of explanation for someone?" (482, original emphasis). There are a few differences between Grimm's question and the one with which we are concerned. One such difference, which is slight, is that we are not especially concerned, as Grimm seems to be, with determining *what it is* in virtue of which

[6] We say "typically", because the lottery could have been rigged, or the winner may not have known that a ticket was bought for her, and so on.

[7] Chanciness is not necessarily the same as randomness. Here is a way of making the distinction: randomness is a feature of a sequence of events, while chanciness is a feature of the process by which an event is caused or brought about. Randomness is roughly a measure of disorderedness, so we could produce a random sequence even with a non-chancy process by, for example choosing the 3156th through 3163rd digits of pi. Conversely, repeated chance processes could (by chance) turn up a non-random sequence. If we flip a coin ten times in a row and it lands heads each time, the sequence is not random, although it was produced by chance. See Eagle (2012).

explanation is needed. Rather, we are primarily concerned with *when* explanation is needed. Perhaps this is only a terminological difference; perhaps it is a real difference, and one cannot give a *complete* account of when explanation is needed without giving an account of what it is in virtue of which explanation is needed — we only mean to emphasize that we shall be satisfied with an account that sorts phenomena into those that need explanation and those that do not, and we will trust that sufficient philosophical illumination will be given by such an account.

A second, slightly more substantial difference between Grimm's question and ours is that he characterizes the putative explanandum as a *situation*, where we have been using the word 'phenomenon'. Initially, Grimm is intentionally ambiguous about what a situation is (see 482, footnote 1), but later characterizes the notion so that a situation encompasses both a *fact* (what we are calling a *phenomenon*) and a *foil*, which is an unrealized possibility that is "live or relevant, relative to the situation in question." (491) If we have understood Grimm's conception of a situation correctly, then there is a confusing circularity in the conception. On his account, it is a fact-and-foil that stands in need of explanation. Given that his account says that explanation is needed in virtue of the existence of a foil, it seems that he builds the answer into the question. In our judgment, it is clearer to keep what Grimm calls a *fact* as that which stands in need of explanation, and let the foil be part of the answer to the question of why that fact needs explanation. Thus, in describing and evaluating Grimm's account, we will continue to conceive of a situation as the fact itself, not the foil. In so doing, we believe that we are not mischaracterizing Grimm's account, but are simply presenting it clearly in a way that accords with common usage.

A third difference between our questions is that Grimm's is explicitly psychological and agent-relative: he is asking what makes a situation need explanation "for someone". Our question does not presume that whether a phenomenon needs explanation is relativized to a person. As it turns out, our account will indeed make the need for explanation relative, but the reference point is a theory or a set of knowledge, which can be something that a person has, but which can also be something shared by a community. Furthermore, as we have noted previously, our account will be normative in the sense that it will allow that there may be differences between when explanation is needed and when an agent *feels* that explanation is needed.

So let us now consider Grimm's account. He does not introduce it all at once, but in his summary, Grimm says that "a situation stands in need of an explanation for someone in virtue of the person's sense that there are various alternative ways the subject of the situation (a system, say, or a substance that constitutes the 'A' in a fact such as A is F) might have been." (493-494)

To assess Grimm's view, we would like to understand more about this 'modal sense' of a foil (an alternative way the subject of the situation might have been) that Grimm identifies as the ground for explanation being needed. Grimm never precisely defines the notion of a foil, preferring to illustrate it with examples. He does say that the modality involved is typically not "broadly logical" or "metaphysical" (489), but is often related to the capacities of an object. This talk of capacities might lead one to think that the relevant modality is nomic or physical possibility. However, such a thought is challenged by Grimm's example of Harry, who was "raised by a community of delusional J.K. Rowling fans" (492) to believe that windows can turn into frogs under the influence of magic. Such a possibility does not seem to be physical or nomic. Grimm could be understood as describing here a modal sense that is relative to an agent's point of view: the transformation of windows into frogs does not fall within the realm of nomic possibility *for us*, but it is at least plausible that it be considered a nomic possibility for Harry. Nonetheless, a different

example of Grimm's clearly shows that he does not have in mind such a relative conception of nomic possibility as determining the need for explanation. In a footnote (490, footnote 12), Grimm suggests that it can make sense to say that the fact that 2 + 3 = 5 is necessary yet needs explanation. It cannot be said to be nomically possible that 2 + 3 = 5 is not necessary, so it cannot be nomic or physical possibility that determines when explanation is demanded.

This last example raises a difficulty in making sense of Grimm's view. It does not seem that there can be "various alternative ways" that 2 + 3 = 5 could have been. Yet, Grimm claims, if one understands that 2 + 3 = 5 is a proposition, then one can see that *qua* proposition, 2 + 3 = 5 might not have been necessary, since not every proposition is necessary. Thus, there are alternatives to the necessity of the proposition, even if it is metaphysically impossible that 2 + 3 not equal 5. But this position is far too permissive. Every subject of a phenomenon can belong to many different kinds, and it will always belong to a sufficiently general kind such that there are alternative ways a thing of *that* kind could have been. For example, *qua* piece of metal, this penny could not have failed to conduct electricity, but *qua* item of money, it could have failed to conduct electricity. If such a permissive standard of possibility were used, Grimm's account would collapse into triviality and all phenomena would stand in need of explanation. So, to be viable, Grimm's account should be supplemented with a characterization of which kinds are relevant for determining whether a subject of a phenomenon could have been different.

However, we will not attempt to fix up Grimm's account along these lines, for it seems to us that it suffers a fatal flaw, no matter what conception of possibility is used and no matter what limitations are placed on the kinds relative to which possibility is assessed. The problem is very similar to the problem described in the previous section that besets the surprise account. In short, the fact that there are alternative ways for a thing to have been is insufficient for the phenomenon in which the thing is involved to need explanation. Ironically, Grimm's own example illustrates this problem. He describes a corn field in which the corn stalks range in height between six and eight feet. He then notes that the fact that a given stalk's height is seven feet "might elicit your curiosity." (489) Indeed, that fact *might* elicit one's curiosity, but it also might not, and to call such curiosity a "need for explanation", as Grimm does, strikes us as too strong. If Grimm and we were walking through the field together, all of us would conceive the same modal possibilities for the corn stalk's height, yet he would have a need to explain it and we, lacking his curiosity, would not. Thus, on this psychological conception of the need for explanation, the sense that there are alternate ways a thing could have been is not sufficient to account for an explanation to be needed.

However, this case might not be seen as definitive. Perhaps we *ought* to feel the need for an explanation in this case. Perhaps Grimm's sense of curiosity is more aligned with his epistemic obligations than is ours. So, one might wonder if Grimm's account could be improved by making it more explicitly normative and less psychological. The corn stalk case would be no counterexample, perhaps, to such an account. However, consider the lottery winner mentioned in the previous section. In this case, any conception of possibility (that allows for there to be possibilities that are not actual) would have to agree that the actual winner might not have won. Nonetheless, there is no need to explain why the winner won. Indeed, people who think an explanation is needed are making a *mistake*. So the existence of alternatives, whether considered psychologically or mind-independently, cannot be sufficient for an explanation to be needed.

Although Grimm's account fails, we think that the examination of his view can yield important lessons. We agree that there is something right in the idea that a network of possibilities

is involved in the determination of explanatory need, but the notion of possibility needs to be explained in more detail. Moreover, Grimm's relativization of possibility to the type or kind to which something belongs also seems a fruitful strategy for pinning down the correct notion of possibility. What our earlier discussion showed is that, since things belong to many different types, we will need a way to sort out the relevant types from the irrelevant ones, or risk triviality. We turn to these tasks and the goal of constructing a better account of explanatory need in the next section.

IV. The Map Account

We call our account *the map account*; why it is so called will be clear presently. The basic ideas of the map account are fairly straightforward. When we ask for explanation of a phenomenon, the asking itself is possible only against the background of a certain understanding of the world (or a part or an aspect of the world). Let us call such an understanding a 'theory' in the sense of a collection of interrelated concepts and beliefs that allows us to make sense of what has happened, predict what will happen, and decide what we should do. It is the theory that provides us with the concepts and beliefs in terms of which we can ask for an explanation, and the phenomenon *needs* explanation only when it does not fit the theory. And it is also the theory that allows us to see a spectrum of acceptable and unacceptable explanations. The theory is analogous to a map of the world (or a part of it) by means of which we can plan our routes, locate where we are, figure out how we got there, and predict what we will see. When we see something that is not supposed to be there according to the map, we ask why it is there — this is analogous to asking for an explanation of a phenomenon. And when we add what we have just seen to the map, we have to do it in such a way that it will fit the rest of the map — this is analogous to giving an explanation of the phenomenon by revising the theory.[8]

In what follows we will spell out these ideas and their implications. We will begin with a rather elaborate example. Imagine you are wandering through the Taï forest in the Ivory Coast. Walking through the dense growth, you are rarely able to see more than 15 feet in front of you, and the calls of insects, birds and other animals provide a constant, cacophonous sonic backdrop. Your path takes you next to a valley, which allows you to see a bit farther, to the opposite side of the valley. As you walk you hear a loud screeching from a group of trees directly across the valley from you. Are those birds? Monkeys? Chimpanzees? Do chimpanzees even spend time in trees? As you continue to walk, the screeching stops, and you find yourself across from a different set of trees, where a new round of screeching begins. Is that the same kind of screeching you heard before? Are those the same particular animals as before, or a new group? Your walk takes you across from a third set of trees, where you hear similar screeching. Once those screeches have died down, you hear no more, and you continue on to camp. You might wonder what kind of animal made that noise. You might even wonder why you heard it from three separate areas, but you will almost certainly not think there is anything about the screeching that *needs* explanation. You will simply take the noise to be part of the general cacophony of the forest.

[8] We are not the first to use the map analogy in the philosophy of science. Toulmin (1953), for example, discusses at great length "the analogy between physical theories and maps", which he thinks "can be used to illuminate some dark and dusty corners in the philosophy of science" (105). A more recent example is Kitcher (2001), which uses the history of map-making to illustrate the notion of accuracy and whether there can be a single ideal theory (see Chapter 5).

As it turns out, you were not alone on your walk. You were accompanied by a primatologist who has spent many years studying the monkeys of the Taï forest. She tells quite a different story about your walk. From her experience, she was instantly able to tell that all three episodes of screeching were produced by different groups of Diana monkeys. When she first learned their calls, she recognized the vocalizations you heard as alarm calls. She earlier discovered that playing pre-recorded sounds of leopards and eagles (both natural predators of Diana monkeys) elicited Diana alarm calls that sounded the same to her ears. However, upon analyzing recordings of the alarm calls using a signal processor, she recognized that there was a small difference between the calls for leopards and the calls for eagles, which she then learned to recognize by ear alone. When, on your walk, you first heard screeching, she recognized the screeching as the eagle alarm call, and looked to the sky to confirm that eagles were flying overhead, which they were. You had not even thought to look up, and had not noticed the eagles at all. She also confirmed that the monkey calls were coming from the lower branches of the trees, where the monkeys go to protect themselves from the eagles. The primatologist recognized the second episode of screeching as also consisting of eagle alarm calls from Diana monkeys, and the eagles were flying in plain view of the location of this second set of monkeys. However the third episode of screeching consisted of leopard alarm calls. The primatologist was puzzled. She had never before heard leopard alarm calls in the presence of eagles. Moreover, the calls seemed to be coming from the lower branches of the trees, which was highly unusual in the case of leopard alarm calls, as the monkeys making such calls usually flee to higher branches to escape the leopards. She wondered if she could have been mistaken about what call she heard. She also wondered if this was a different sort of Diana monkey, or different sort of monkey altogether that was making the calls. However, in her experience, the alarm calls of all the primates in the forest were recognizable by all the other species, so it would be very unusual if a different sort of monkey used the leopard call to warn of eagles. Of course, she also wondered if there was a leopard nearby, although the fact that the monkeys did not flee to the higher branches seemed inconsistent with this possibility. With a slightly brisker step, the primatologist returned to camp with a host of new questions, new ideas for further research, and the recognition that the phenomenon she had just observed *needed explanation*.[9]

There are a number of observations about this story that are relevant to our investigation: (a) The third episode of monkey screeching did not need an explanation from your point of view, but it did from the point of view of the primatologist. This was not because the primatologist was surprised by the phenomenon, while you were not — you *were* surprised by it, yet you did not think the phenomenon needed an explanation (so the surprise account is inadequate). Nor was it because the primatologist thought of an alternative to the phenomenon, while you did not — you had been entertaining the idea of a much quieter forest, but you still did not think the phenomenon needed an explanation (so the fact-and-foil account is inadequate). Part of the difference between you and the primatologist is that you did not even recognize the phenomenon as perceptually distinct from other relevant phenomena, other than that it occurred in a different place at a different time. But even had you been able to distinguish the sounds, you had no background understanding that would have told you that the difference was salient.

[9] This example is loosely drawn from a story about the work of Klaus Zuberbühler as reported on the Radiolab program "Wild Talk". It has been tailored for narrative purposes, and is not necessarily an accurate description of phenomena of the Taï Forest.

(b) The need for explanation arose for the primatologist as a result of a conflict between her background understanding or theory and the phenomenon. Her theory about the monkeys of the Taï forest was, as it were, her map of the forest — not a complete map of the forest, but at least a map of the forest as far as the monkeys are concerned. For her, the third episode of monkey calls needed an explanation because it did not fit her map. According to her map, either the eagles should not have been there, or the leopard alarm calls should not have been there. It was relative to her map that the leopard alarms calls in the presence of eagles required an explanation. Notice also that it was her map that instructed her to look for and observe relevant aspects of the situation that might not be evident from her initial perception, such as whether there were eagles flying nearby and whether the calls came from the lower or higher branches of the trees.

(c) The difference between the primatologist and you is not that she had a map of the forest and you did not. Given that you had some understanding of what a forest is and what kinds of things are supposed to be in a forest, you had a map too. It's just that your map was much more coarse-grained than the primatologist's. Your map, for example, indicated that there are monkeys in the forest, but it did not distinguish between different kinds of monkeys there are, let alone tell you the behavior of Diana monkeys. Your map probably told you monkeys sometimes screech. If this was all it told you about monkey screeching, then all three episodes of monkey screeching fit your map — none of them needed an explanation. Although your map did not tell you that the screeching needed explanation, it would imply that other phenomena need explanation. Imagine that you saw a leopard being chased by a monkey. This phenomenon would not fit your map, for your map tells you leopards are predators of monkeys, not the other way around. You would think the phenomenon needed an explanation.

(d) The need for explanation in this case was not purely theoretical: one reason why the phenomenon needed an explanation was that one aspect of some potential explanations, the presence of a leopard nearby (besides the eagles), had practical consequences for the defenseless primatologist walking through the forest. The map against which the phenomenon needed an explanation was also a map she actually used to guide herself through the forest. She had been using that map up to that point and the map had been correct and useful. After observing the third episode of monkey calls, however, she could not trust the map the way she had before. She had to either revise the map or find a way of resolving the apparent conflict between the current map and the phenomenon in question.

These observations can be generalized and augmented:

(A) It may appear that the need for explanation is relativized to a person, that is, the same phenomenon may need explanation to one person but not to another. However, since it is the map a person uses that determines whether the phenomenon needs explanation, and a map can be shared by a group of people or a community (such as a scientific community), the need for explanation is not fundamentally agent-relative. In fact, even for the same person who thinks the phenomenon needs explanation, she might not have thought so had she seen the same phenomenon at a different time, for she might not have used the map relative to which the need for explanation arises. Suppose the primatologist in our example has an assistant who has just started being trained by her to observe the special behavior of the monkeys of the Taï forest. Since the assistant is still a novice, he will sometimes miss important signs that are obvious to the primatologist. Now suppose in our story the primatologist was accompanied by the assistant. During the third episode of screeching, which was puzzling to her, she noticed that the assistant did not seem to be puzzled at all. She

decided not to tell him directly what puzzled her; instead, she simply pointed to the eagles flying overhead. The assistant noticed immediately that the presence of the eagles did not match the leopard alarm calls, which he did recognize, and said, "That's strange!" He thought the phenomenon needed an explanation, but he probably would not have thought so had he been on his own when this happened. Using the map analogy, we may say that he had a similar map as the primatologist's, but he had not mastered how to use it yet.

(B) The need for explanation arises when there is a conflict between the map used and a phenomenon that is within the mapped area, that is, when the phenomenon does not fit the map. There are three main ways for a phenomenon P to not fit a map M: (i) P is not on M because P is either incompatible with the rest of M or highly unlikely given the rest of M — if you are guided by M, you expect to *not* see P; (ii) P is different from what M indicates — if you are guided by M, you do expect to see something like P, but not quite the way P actually is; (iii) P consists in the absence of something X that is supposed to be there according to M — if you are guided by M, you expect to see X, but you see P instead (P may just be the absence of X or a phenomenon that is utterly different from X).[10] Sometimes it is, however, not clear which of these three kinds of cases a phenomenon belongs to so that it requires an explanation, for the same phenomenon can be described in more than one way even given the same map. Indeed, the third episode of screeching in our example could be understood as a case of (i) or (iii). If the phenomenon is described as "leopard alarm calls in the presence of eagles", then it was a case of (i); if the phenomenon is described as "no eagle alarms calls in the presence of eagles", then it was a case of (iii). But as long as P is understood as a case of (i), (ii), or (iii), it does not fit the map and needs explanation.

(C) Different people (or groups of people) can have different maps of the same area. Some of these maps are incompatible so that at least some of them are incorrect representations (incorrect in at least some respects) of the area. Although two different maps of the same area can both be correct if they are about different aspects of the area (such as a biological map versus a cultural map), it is rare for two maps of the same area to be totally unrelated in such a way that there *cannot* be incompatibility between them no matter how they are revised. Another way for two different maps of the same area to be both correct is that one of them (M1) is less fine-grained than the other (M2). M1 and M2 are both correct in the sense that there is nothing incorrect on them, but M2 has all the correct details that M1 has *and more*. If a phenomenon P does not fit M1, it is impossible for it to fit M2; but if P does not fit M2, it may still fit M1. And when P does not fit M2, a user of M2 has to consider the possibility that M2 is incorrect as far as P is concerned.

(D) Usually a person thinks a phenomenon needs explanation not simply because the phenomenon does not fit *a map she has*, but because it does not fit *the map she is using*. The difference between having a map (without using it) and using a map is this: having a map is accepting it as an accurate representation of the area, while using a map is accepting it as an accurate representation of the area *and* being guided by it through the area. M1 and M2 can both be correct representations of area R, though they represent different aspects (or sets of aspects) of R. The primatologist in our

[10] In (ii) and (iii), as in (i), P is either incompatible with the rest of M or highly unlikely given the rest of M. It should be noted that it is not clear how unlikely P has to be for it to not fit M. This is not a problem for our account because it implies only that there are cases in which it is not clear whether the phenomenon needs explanation. We would like to thank an anonymous reviewer for pointing out the difference between P's being incompatible with the rest of M and P's being highly unlikely given the rest of M.

example was being guided by her map of the forest as far as the monkeys there are concerned, but she has another map of the forest that was not being used when she was walking through the forest, namely, a map of the forest as far as the archaeological sites there are concerned. She has some interest in archaeology, has read quite a number of books about the archaeological sites in the forest, but when she walked through the forest she was doing primatological research, not archaeological research (nor will she be doing any archaeological research in the future). So even if she saw something that did not seem to fit her archaeological map of the forest, she might not think it needed an explanation — she might simply ignore it and move on because she had only limited time, energy, and resources.

Our account seems to suggest that any description of a phenomenon is composed of the elements of a certain map, if not actually part of the map. It is true that any description of a phenomenon is part of a certain collection of interrelated concepts and beliefs, but not every such collection is a map in our sense. A collection of interrelated concepts and beliefs is a map if and only if by means of it we can make sense of what has happened, predict what will happen, and decide what we should do. This helps explain why a phenomenon that is highly unlikely may not need explanation — the phenomenon is highly unlikely only under a certain description, while the description is not composed of the elements of any particular map.

Consider two scenarios, both concerning the temporal evolution of a system consisting of a glass of water in which a single ice cube has been placed. In the first scenario, after ten minutes have passed, the ice cube has completely melted. In the second, after the same ten minutes, the entire contents of the glass are frozen solid. Surely, under normal circumstances, we would think that the phenomenon in the second scenario needs explanation, but that in the first does not. It is tempting to think that this is simply because in the first scenario the final state is likely, while in the second scenario the final state is highly unlikely. But note that if we describe the final state of the system in the two scenarios in terms of the microscopic description of the state of all the component particles, then each final state is equally unlikely relative to the initial condition of the system.[11] Under the microscopic description, the final state of the system in the first scenario is highly unlikely, but it still does not need an explanation. According to the map account, this is because the microscopic description is not composed of the elements of any particular map — there is no map by means of which we can understand why the temporal evolution of the kind of system in question results in a certain final state rather than another at the microscopic level, by means of which we can predict what will happen, and by means of which we can decide what we should do.

The same analysis is applicable to the lottery example we discussed in section II. The lottery winner is surprised that she won the lottery because the probability of winning was so low. Indeed, the probability of winning was low no matter who won the lottery. Call the winner *Winnie*. The fact that Winnie won the lottery does not need an explanation because the description 'Winnie won the lottery' is not composed of the elements of any particular map by means of which we can understand why a particular person (rather than another person) won the lottery, by means of

[11] Or equally likely, if one insists on describing the initial state in terms of the microscopic description and the systems evolve deterministically. For the sake of argument let us assume that under the microscopic descriptions both final states are highly unlikely. The example does not, of course, suggest that it is as likely that the water will freeze as that it will stay liquid if the temperature in the room does not change — we are speaking only of the two final states *under the microscopic description*.

which we can predict who will win the lottery next time, and by means of which we can decide what we should do as far as playing the lottery is concerned. The same is true of the descriptions 'A woman won the lottery', 'My neighbor won the lottery' (if Winnie is my neighbor), 'A woman who owns a Ford won the lottery' (if Winnie owns a Ford), and many others.

How about the example of the colleague who won the lottery three times in a row and every time it was exactly one day after he received a fortune cookie saying that he would win the lottery the next day? Call this colleague *Winfred*. We would think that Winfred's winning this way needs explanation. According to the map account, this is because there is a map we are using such that Winfred's winning this way does not fit the map. What is the map? It cannot be the map about lotteries *only* because such a map has nothing to do with why and how particular individuals won the lottery. Can it be the map about cookies? Most of us do have beliefs about cookies, but those beliefs do not necessarily form a map. A baker presumably has a map about cookies, but such a map has nothing to do with lotteries.

It should now be clear that a phenomenon requires explanation only under a certain description, and that very description has to be in terms of the elements of the map being used which the phenomenon does not fit. This is why the map in Winfred's case cannot be a map about cookies, at least not one that a baker has. So, what map are we using when we think Winfred's winning requires explanation? A reasonable suggestion is that it is our map about predictions: according to this map, fortune-telling, including fortune-telling by fortune-cookies, is generally unreliable because a prediction is reliable only if it is based on some well-established causal understanding of how the things involved work. It is because Winfred's winning does not fit this map that we think it needs explanation.[12]

If Winfred won only once this way, we could reasonably say it was a coincidence; but three times in a row is probably too many times to dismiss as coincidences. Here people may have different intuitions, but if three times in a row is not sufficient for people to think an explanation is needed, we can enlarge the number; sooner or later people would think an explanation is needed, and the map account would still give the same analysis of why an explanation is needed — according to the map used, it should not have happened.

Some may think the map analogy is no different from the well-known analogy of a 'web of belief'.[13] This is not the case. While the web analogy is usually used to represent our whole belief system, the map analogy as we are using it here can be used to represent both our whole belief system and any of the sub-systems we have. There is only one big web, but there can be many maps. Besides, the map analogy allows for the possibility of our different smaller maps not forming one big map — they simply don't form a coherent whole and some of these maps may even be incompatible. These maps may give us what Nancy Cartwright (1999) calls a "dappled world". We are not suggesting that there cannot be a coherent belief system about the world as a whole, but our

[12] We can, however, imagine that Winfred himself does not think his winning three times in a row requires explanation if we imagine that he is using a map very different from ours: his map includes reliable fortune-telling as part of reality and includes details of how certain fortune-telling objects, such as fortune cookies, work. And the reason why Winfred does not think his winning needs explanation is simply that he believes there is already an explanation — the fortune cookies had correctly predicted that he would win.

[13] The most famous passage is probably Quine (1951: 42-43), though Quine uses the word 'fabric' instead of the word 'web'. See also Quine & Ullian (1978).

map account at least does not presuppose that there can be such a belief system. Another difference between the map analogy and the web analogy is that the former allows for the possibility of there being gaps or missing parts, while the latter does not seem to allow for that. This is exactly why the map analogy is more suitable for giving an account of the need for explanation (an explanation is needed when there is a gap or missing part).

In section IV (D) we made the distinction between a person's merely having a map and her also using the map: a person can have a map without using it, that is, having the understanding involved without applying it to anything or being guided by it. As will be seen in the next section, this distinction is important to understanding the normative force of the judgment that a phenomenon needs explanation. This distinction between having and using also marks a difference between the map analogy and the web analogy, for the distinction cannot be made in the web analogy.[14]

V. Why the Map Account Is Normative

The above account of the need for explanation relies on the map analogy, and as Stephen Toulmin so crisply puts it, "like any analogy, it will take us only a certain way." (Toulmin 1953: 105). The analogy has not, however, taken us far enough, for we still have to make use of it to explain why our account is a normative one.

The map account, as we have characterized it so far, is clearly not a psychological account, for whether a phenomenon fits a particular map, which determines whether the phenomenon needs explanation, is not a psychological issue (though its fitting or not fitting the map may be accompanied by some psychological state). We need to, however, make clear why it is a normative account.

As we explained in section I, an account of the need for explanation is normative if it specifies the need for explanation in terms of there being a *good reason* for us to look for an explanation — if a phenomenon is such that there is a good reason for us to look for an explanation of it, then the phenomenon needs an explanation; and if the phenomenon needs an explanation in this way, then we *should* look for an explanation of it no matter how we feel when we are being faced with the phenomenon. 'Having a good reason to look for an explanation' is a directive, for it directs or requires us to do something, namely, look for an explanation. However, whether we are subject to such a directive depends not only on whether the phenomenon involved fits the map, but also on whether we are *using* the map rather than merely having it. If S merely has the map M, then even though her seeking an explanation of P (where P does not fit M) would not be considered unjustified, she is not directed or required to do so. And in cases in which the using of M is an essential part of an inquiry, S is epistemically obliged to look for an explanation of P, for not doing so conflicts with the purpose of inquiry.

In general, when a phenomenon does not fit the map we are using, we have good reason to look for an explanation of it, and the request for an explanation is in this sense justified. But there

[14] We would like to thank an anonymous reviewer for urging us to clarify the difference between the map analogy and the web analogy.

are at least two ways that such a request for explanation can fail to be justified. The first way is that the map being used as the basis for the request fails to be justified. Call this the *wrong map error.* Consider, for example, a proponent of the Flat Earth Theory (FE). According to this theory, the earth is a flat disc moving with constant acceleration in the "up" direction. Our advocate (an FE'er) recognizes that any water on such a disc not contained in a depression would flow off the sides, yet knows that the ocean level has not been constantly diminishing (indeed, it is increasing). So for the FE'er, the phenomenon that the ocean level is not diminishing may seem to need an explanation (which might be that there is a ring of mountains around the edge of the disc keeping the oceans trapped ('FAQ' 2012)). However, the particular map (both literal and figurative) that the FE'er uses is not well supported by evidence, and thus the phenomenon does not in fact need an explanation.

A second way in which the demand for explanation may fail to be justified is that one may incorrectly judge that a phenomenon does not fit a map. Call this the *no false fit error.* Consider a student learning simple Mendelian genetics who mistakenly calculates that the offspring of two heterozygous parents should be one third homozygous dominant, one third homozygous recessive, and one third heterozygous (rather than one quarter of each homozygous pair and one half heterozygous). The student may well judge that her classroom botanical experiments fail to fit her map. But of course, no explanation is needed, as the student has simply failed to draw the correct inferences from the map. In general, not every such failure to recognize the implications of a map should count as failing to recognize that a phenomenon fits the map. If such were the case, then our general lack of logical and computational omniscience would mean that many phenomena which we normally think need explanation would not in fact need explanation. For example, the possibility of black holes is a consequence of the theory of relativity, but it took concerted mathematical work to show that this was the case. Counterfactually from what actually happened, had black holes been observed prior to this work being done, it would have been *reasonable* to say that the phenomenon of black holes did not fit the map provided by general relativity, and thus would need an explanation, although their possibility is implied by the theory. Thus, we should not identify a map with the logical closure of a theory, as we want to allow that the logical consequences of some theory do not fit with the map provided by the theory at the time. We cannot easily say when the consequence of a theory belongs to the map associated with the theory, and when it doesn't; we can only say that the consequence belongs to the map when its adherents *should have known* it.

Note that these two different ways in which a demand for explanation may not be justified are roughly analogous to two ways in which an argument may be unsound: the wrong map error is analogous to the falsity of the premises, and the no false fit error is analogous to the invalidity of the argument. Arguments are typically enthymematic, and in charitably assessing a bad argument one often has to supply missing premises. Sometimes, one must choose between providing false premises that make the argument valid, or true premises that make it invalid. In such cases there is no fact about which way the argument fails. Similarly, it may be the case with many phenomena that fail to need explanation that the problem could be either of our two errors, and there is no fact about which is the right diagnosis.

Before we assess our account in the conclusion, there is one more aspect of our account worth mentioning. At the beginning of this investigation, we did not explicitly presume a theory of explanation. Indeed, we expressed the hope that an account of explanatory need might shed some light on the nature of explanation. However, it is hard not to see the affinities between our account

and the view of explanation as unification. According to that view (see, for example, Kitcher 1981) an explanation is a derivation or argument drawn from a set of derivations that unify our knowledge. This set, called an *explanatory store*, is more unifying the more stringent the argument patterns it contains and the greater the logical strength of the set. The strength of an explanation is a function of how unifying is the explanatory store from which it is drawn. It seems reasonable to roughly identify the explanatory store with a map, or at least to recognize that the store derives from the map. If a phenomenon does not fit on the map, then the unifying power of the map is diminished in two ways. As with any unexplained phenomenon, a phenomenon that does not fit on the map is a fact that fails to be unified by the explanatory store. But unlike just any unexplained phenomenon, the phenomenon that does not fit on the map also diminishes the stringency of the explanatory store. If the phenomenon does not fit, then that is because an argument pattern from the store implies that the phenomenon should not exist, or it should be different from how it is. Thus, to accommodate the phenomenon, the argument pattern must be augmented with provisos that make the pattern more complex and less stringent. To the degree that logical strength and stringency of one's explanatory store is threatened by a phenomenon, one should aim to accommodate that phenomenon. So it is natural to see that an advocate of explanation as unification would find our account of the need for explanation motivated by her view.

Although the unification view of explanation may provide a motivation for our account, the account was primarily developed through reflection on other views of explanatory need and specific cases. Nonetheless, at root there is a similar intuition underlying both views, which is to recognize how explanation fits phenomena into broader patterns. However, this intuition is arguably at the root of other views of explanation as well, and our account does not imply the unification view of explanation. Indeed, one may see our imprecise characterization of a map to be a virtue because it allows our account to be tailored to fit more than one view of explanation. For the covering law model of explanation, the map would consist of laws of nature, some initial conditions and some logical consequences of the laws and initial conditions. On a causal view of explanation, the map would consist of a causal structure. And for a pluralist about explanation, different kinds of maps would be appropriate for different circumstances.

VI. Concluding Remarks

The starting point of our investigation is the fact that we don't think *every* heretofore unexplained phenomenon needs explanation. We seem to have the understanding that some (heretofore unexplained) phenomena need explanation and some do not, and our practice reflects such understanding — when something has happened and we don't know why it happened, sometimes we will demand an explanation and sometimes we won't. There seems to be a distinction between phenomena that need explanation and phenomena that do not, and the map account is an attempt to draw the distinction *in the right way*. In other words, we have aimed to provide an account that captures our pre-theoretical understanding of the need for explanation, that is in accordance with our practice of demanding an explanation, and that sheds light on how the need for explanation is related to other important aspects of explanation. We believe we have succeeded in all three respects.

Our pre-theoretical understanding of the need for explanation seems to be that there is a *real* distinction between phenomena that need explanation and phenomena that do not: whether a

phenomenon needs explanation is not merely a matter of whether we happen to want to have an explanation of it. This independence of the need for explanation is related to another aspect of our pre-theoretical understanding of the need for explanation: if a phenomenon needs explanation, we should be able to point out what it is about the phenomenon that makes it require an explanation. If the need for explanation is not identical with our desire for explanation, then there should be something about the phenomenon (rather than about us) that accounts for this difference.

A psychological account may find a way to capture such pre-theoretical understanding of the need for explanation, but it is clear that the map account, which is normative, achieves this aim straightforwardly. According to the map account, there is a real distinction between phenomena that need explanation and those that do not: a phenomenon needs explanation when it does not fit the map we are using, and whether it fits the map is *a matter of fact* independent of our assessment. We realize that the phenomenon needs explanation when we discover that it does not fit the map. And since it is relative to the map that the phenomenon needs explanation, we usually can easily point out what it is about the phenomenon that makes it need explanation — we can simply point out the way it does not fit the map.

Although the account we have proposed is normative, we have not assumed that the normative and the psychological are unrelated. As we remarked in section II after criticizing the surprise account, the surprisingness of a phenomenon, when properly qualified, is indeed relevant to the phenomenon's need for explanation. Given the map account, the surprisingness that is relevant to the need for explanation is not a brute psychological state that we happen to be in; it is a form of *intellectual* curiosity that is not possible without the context of map using. Strictly speaking we are surprised not by the phenomenon, but by the phenomenon's not fitting the map used. If we did not have such intellectual curiosity about the phenomenon, then even if we were still surprised by it for some other reason, we would not think the phenomenon needs explanation.

Take the lottery cases again. Although Winnie is surprised that she won the lottery, she should not think her winning requires an explanation. Winnie is surprised that she won simply because the probability of winning a lottery is extremely low, but her winning does not need an explanation because it was not in conflict with any relevant map she is (or we are) using; it fits all her (or our) maps. Indeed, as we have seen, there can be descriptions of her winning that do not belong to any map. By contrast, Winfred's winning the lottery three times in a row is not just surprising; it requires an explanation because it does not fit at least one map that we are using and we are intellectually curious because of that.

Both lottery cases are examples of how the map account matches our practice of demanding an explanation. It would take empirical research to establish that in most, if not all, cases in which we do (or do not) request an explanation of a phenomenon, the map account would say we should (or should not) demand one. However, in light of the lottery cases and similar examples, it is at least reasonable to claim that the map account is in accordance with our practice.

If the map account is correct, it is easy to see why and how, as many philosophers of science implicitly assume in their discussions of explanation, scientific explanation and everyday explanation are similar in structure. The map account is supposed to apply to both kinds of explanation. Scientific explanation and everyday explanation are similar in structure because in both cases the need for explanation arises when a phenomenon does not fit the map used. In the case of scientific explanation, the map is a scientific theory; in the case of everyday explanation, the

map is our everyday understanding of the world (or a part or an aspect of the world), which can be loosely called 'a theory' (we have everyday theories about physical objects, animals and plants, seasons, weather, human motivations, etc.). If both the need for scientific explanation and the need for everyday explanation arise when a phenomenon does not fit the map used, then in both cases the explanation given is either an attempt to fix the map or an attempt to show that the phenomenon was not really in conflict with the map.

Given that explanation is an essential part of science, it is reasonable to assume that the need for explanation is one of the determining factors of the development and progress of science. While neither the surprise account nor the fact-and-foil account can help us understand how exactly the need for explanation is related to the development and progress of science, the map account gives us a clear picture of how the two are related: since the need for explanation of a phenomenon arises only when the phenomenon does not fit the map used, the need for explanation usually results in an improvement of the map, improvement in the sense that it fits what it is a map of better than before. The improvement of a scientific map (i.e. a scientific theory) is the development and progress of a branch of science. Since scientific maps are interrelated, the improvement of a scientific map can affect other scientific maps and may result in an improvement of them too. So, it is not an exaggeration to say that the need for explanation is (partially) what pushes science forward.

Acknowledgements We are grateful to two anonymous reviewers for helpful comments.

References

Cartwright, N. (1999). *The Dappled World*. Cambridge: Cambridge University Press.

Eagle, A. (2012). 'Chance versus Randomness', *The Stanford Encyclopedia of Philosophy (Spring 2012 Edition)*, E. N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2012/entries/chance-randomness/>.

'FAQ' (2012, August 19). New Flat Earth Wiki. Retrieved from http://www.theflatearthsociety.org/wiki/index.php?title=FAQ

Grimm, S. R. (2008). 'Explanatory Inquiry and the Need for Explanation', *British Journal for the Philosophy of Science*, 59: 481-497.

Hempel, K. (1965). *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York: Free Press.

Kitcher, P. (1981). 'Explanatory Unification', *Philosophy of Science*, 48: 507-31.

Kitcher, P. (2001). *Science, Truth, and Democracy*. Oxford: Oxford University Press.

Korsgaard, C. M. (1996). *The Sources of Normativity*. Cambridge: Cambridge University Press.

Laudan, L. (1977). *Progress and Its Problems: Towards a Theory of Scientific Growth*. Berkeley: University of California Press.

Quine, W. V. (1951). 'Two Dogmas of Empiricism', reprinted in his *From a Logical Point of View*, 2nd edn. Cambridge: Harvard University Press, 20–46.

Quine, W.V. & Ullian, J. S. (1978). *The Web of Belief*, 2nd edn. New York: Random House.

Salmon, W. (1978). 'Why ask, "Why?"? An Inquiry concerning Scientific Explanation', *Proceedings and addresses of the American Philosophical Association*, 51: 683-705.

Thomson, J. J. (2008). *Normativity*. Chicago and La Salle: Open Court.

Toulmin, S. (1953). *Philosophy of Science: An Introduction*. New York: Harper and Row.

Woodward, J. (2009). 'Scientific Explanation', *The Stanford Encyclopedia of Philosophy (Winter 2011 Edition)*, E. N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2011/entries/scientific-explanation/>.